



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/813,843	03/31/2004	Nitzan Peleg	200311278-1	6064
22879 7590 12/31/2007 HEWLETT PACKARD COMPANY P O BOX 272400, 3404 E. HARMONY ROAD INTELLECTUAL PROPERTY ADMINISTRATION FORT COLLINS, CO 80527-2400			EXAMINER LIE, ANGELA M	
			ART UNIT	PAPER NUMBER
			2163	
			NOTIFICATION DATE	DELIVERY MODE
			12/31/2007	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

JERRY.SHORMA@HP.COM
mkraft@hp.com
ipa.mail@hp.com



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

MAILED

DEC 26 2007

Technology Center 2100

**BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES**

Application Number: 10/813,843
Filing Date: March 31, 2004
Appellant(s): PELEG ET AL.

Barry D. Blount
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed August 27, 2007 appealing from the Office action mailed March 31, 2007.

(1) Real Party in Interest

A statement identifying by name the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of Claims

The statement of the status of claims contained in the brief is correct.

(4) Status of Amendments After Final

The appellant's statement of the status of amendments after final rejection contained in the brief is correct.

(5) Summary of Claimed Subject Matter

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

The appellant's statement of the grounds of rejection to be reviewed on appeal is correct.

(7) Claims Appendix

The copy of the appealed claims contained in the Appendix to the brief is correct.

(8) Evidence Relied Upon

5963959	Sun et al	10-1999
7111020	Gupta et al	3-2002

(9) Grounds of Rejection

The following ground(s) of rejection are applicable to the appealed claims:

Claim Rejections – 35 U.S.C – 102

6. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

7. Claims 1, 3, 5, 11-17, and 20 are rejected under 35 U.S.C. 102(b) as being anticipated by Sun et al (US Patent No. 5, 963,959, Date of Patent: October 5, 1999).

Claim 1:

Regarding Claim 1, Sun teaches a system for performing refresh operations, the system comprising:

a base table having a first plurality of data entries (Figures 2A,B, C, diagram 200, Sun);

a first materialized view that comprises a second plurality of data entries, the second plurality of data entries being associated with the first plurality of data entries in the base table (Figures 2A,B, all features, wherein defined in column 4, lines 29-41, wherein a series of modifications a user might make to a master table and the corresponding entries recorded in a master log and master table 200 within FIG. 2(a) is a table of customer information including a column for a primary key CID, a customer identifier, and a column ZIP for a customer's ZIP code, wherein each row represents a particular customer, who is assigned a non-null, unique identifier, CID, wherein the

corresponding master log 210 is empty and wherein master table 200 of FIG. 2(b) is the result of adding a new customer with a CID of 5 and a ZIP of 22046 to master table 200 of FIG. 2(a), wherein the primary key value of the inserted row, 5, is recorded in master log 210, Sun);

a refresh log that contains a plurality of changes in the base table (Figure 2C, wherein column 4, lines 42-49, the result of deleting the customer identified CID of 2 from the master table, the primary key value of 2 is stored as a new entry in master log, wherein if the zip code of customer CID of 4 in master table is changed from 22090 to 20190, then the master table is the result, the primary key value 4 of the updated row is stored as a new entry in master log, Sun); and

a module adapted to perform a refresh operation on the first materialized view using the second plurality of data entries, the module configured to (Figure 3, all features, wherein column 5, lines 10-15, the operation of a fast refresh mechanism, wherein the primary key values are selected from the master log which are not found in the master view the, the result of reissuing the snapshot definition query on the master table, Sun);

access the refresh log and the first materialized view (column 6, lines 65-66, wherein master table is accessed by the primary key values recorded in the master log, Sun);

calculate a plurality of delta values from the plurality of changes in the refresh log and the second plurality of data entries in the first materialized view (column 6, lines 43-45, wherein two new rows with column primary key, i.e. CID, of 5 and 6 are added, resulting in snapshot, diagram 400 within Figure 4e, Sun);

apply the plurality of delta values to the second plurality of data entries in the first materialized view (Figures 7A and B, all features, wherein defined in column 8, lines 52-67, Sun); and

provide the plurality of delta values to a delta adaptation module for updating a second materialized view (column 9, lines 27-49, Sun).

Claim 3:

Regarding Claim 3, Sun teaches wherein the first plurality of data entries and the second plurality of data entries each include one of a plurality of grouping identifiers that associate each of the first plurality of data entries with the second plurality of data entries (Figures 2A,B, and C, all features, wherein 2A, illustrates a plurality of data entries, and Figure 2B, illustrates another second plurality of data entries, and Figure 2C, diagram 210, illustrates a plurality of CID, i.e. column for primary key, which is equivalent to group identifiers, which is associated with Figures 2A, and B, which are the first and second data entries, Sun).

Claim 5:

Regarding Claim 5, Sun teaches wherein the second plurality of data entries each comprises a grouping field and a count field (Figure 6C, contains a CID, equivalent to a group field and Time\$\$ which is equivalent to count field, wherein it's a value that is added, wherein clarified in column 7, lines 51-58, Sun).

Claim 11:

Regarding Claim 11, Sun teaches a system for performing a refresh operation, comprising:

means for deriving a first materialized view from at least one base table (Refer to claim 1, wherein this limitation substantially the same/or similar, Sun);

means for accessing a refresh log and the first materialized view to perform the refresh operation on the first materialized view (Refer to claim 1, wherein this limitation substantially the same/or similar, Sun);

means for calculating a plurality of delta values by combining a plurality of changes in the refresh log and a plurality of entries in the first materialized view (Refer to claim 1, wherein this limitation substantially the same/or similar, Sun);

means for applying the plurality of delta values to the first materialized view (Refer to claim 1, wherein this limitation substantially the same/or similar, Sun); and

means for providing the plurality of delta values to a delta adaptation module for refreshing a second materialized view (Refer to claim 1, wherein this limitation substantially the same/or similar, Sun).

Claim 12:

Regarding Claim 12, Sun teaches a method of performing a refresh operation, the method comprising:

deriving a first materialized view from a base table (Refer to claim 1, wherein this limitation substantially the same/or similar, Sun);

obtaining a refresh log and the first materialized view to perform the refresh operation on the first materialized view (Refer to claim 1, wherein this limitation substantially the same/or similar, Sun);

calculating a plurality of delta values by combining a plurality of changes in the

refresh log and a plurality of entries in the first materialized view (Refer to claim 1, wherein this limitation substantially the same/or similar, Sun);

applying the plurality of delta values to the first materialized view (Refer to claim 1, wherein this limitation substantially the same/or similar, Sun); and

providing the plurality of delta values to a delta adaptation module for refreshing a second materialized view derived from the first materialized view (Refer to claim 1, wherein this limitation substantially the same/or similar, Sun).

Claim 13:

Regarding Claim 13, Sun teaches wherein obtaining and calculating are performed in a database management system ("DBMS") (column 5, lines 6-9, Sun).

Claim 14:

Regarding Claim 14, Sun teaches wherein applying the plurality of delta values comprises utilizing a plurality of operators to modify the first materialized view (column 8, lines 58-60, Sun).

Claim 15:

Regarding Claim 15, Sun teaches wherein the plurality of delta values to a delta processing module that applies the plurality of delta values to the first materialized view (see abstract, wherein the primary key values of the modified rows are recorded in a master log, wherein in response to a fresh command differences between the master table and snapshot are reconciled based on primary key values stored in the master table, the master log, and the snapshot, Sun).

Claim 16:

Regarding Claim 16, Sun teaches wherein processing the plurality of delta values in the delta adaptation module to create a plurality of second materialized view changes for

the second materialized view (Figures 7A and 7B, wherein diagram 710 illustrates updateable snapshot log, and wherein 7B, diagram 710 illustrates snapshot log which is equivalent to a delta adaptation module, wherein OLD\$\$ indicates a primary key value is old or new, MOD\$\$, indicates insert, delete, update, and TIME\$\$ illustrates a refresh timestamp, and CID and TIME\$\$, are already created, see diagram 210, Figures 2A,B, and C, in which CID is equivalent to the group identifier; column 8, lines 65-67, wherein updateable snapshot of Figure 7b is the result of adding a new customer with a CID of 5 and a ZIP of 22046 to updateable snapshot, which is equivalent to creating a plurality of second materialized view changes, Sun);

calculating a plurality of second materialized view delta values that represent the plurality of second materialized view changes to be applied to the second materialized view (column 9, lines 1-6, Sun); and

applying the plurality of second materialized view changes to the second materialized view (Figure 7C, all features, wherein MOD\$\$ illustrates insert, i.e. I, and delete, i.e. D, with CID's 5 and 2, in which it is associated with diagram 700 within 7C, wherein CID 2 is deleted from diagram 700 in Figure 7C and wherein Figure 7B, CID 5 is requested to be inserted by the MOD\$\$, wherein it is shown in Figure 7C, that it was inserted, in which a change was made, and viewed. Sun).

Claim 17:

Regarding Claim 17, Sun teaches wherein combining a tuple table with the plurality of delta values and projecting the plurality of second materialized view changes based upon the tuple table and the plurality of delta values (column 4, lines 8-17, wherein a primary key is a set of columns in a table having a combined value that is unique and non-null

within a table, wherein a primary key value is able to uniquely identify each row in the table, and wherein since rows are uniquely identified by primary key values the fast refresh mechanism, employs primary keys, wherein the primary key values of modified rows of a master table are recorded in a master log, Sun).

Claim 20:

Regarding Claim 20, Sun teaches a computer program, comprising:

a machine-readable medium (column 3, lines 59-65, Sun);

a refresh log stored on the machine readable medium, the refresh log containing a plurality of change entries (Refer to claim 1, wherein this limitation is substantially the same/or similar, Sun); and

a refresh manager stored on the machine readable medium, the refresh manager being adapted to refresh a first materialized view derived at least in part from a base table by computing a plurality of delta values in a delta calculation module based on the refresh log (column 8, lines 55-64, wherein MOD\$\$ has three values: 'I' for insert, 'D' for delete, and 'U' for update and the updateable snapshot log 710 has an old/new column, wherein the OLD\$\$, which indicates whether a primary key value for the row is old, i.e., O, or new, i.e., U, or unchanged, i.e., U, which is interpreted to be the "plurality of data values", column 6, lines 43-44, wherein two new rows with CID of 5 and 6 are added resulting in a snapshot, in which this is interpreted to be equivalent to "the refresh manager being adapted to refresh a first materialized view derived at least in part from a base table by computing a plurality of delta values in a delta calculation module based on the refresh log", wherein the refresh manager is interpreted to be the "snapshot", Sun), and the first materialized view, applying the plurality of delta values in a delta processing module to the first materialized

view (Refer to claim 1, wherein this limitation is substantially the same/or similar, Sun), and providing the plurality of delta values to a delta adaptation module derived from the first materialized view (Refer to claim 1, wherein this limitation is substantially the same/or similar, Sun).

Claim Rejections 35 U.S.C – 103

8. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

9. Claims 2, 4, 6-10, 18-19 and 21-22 has been rejected under 35 U.S.C. 103(a) as being unpatentable over anticipated by Sun et al. (US Patent No. 5,963,959, Date of Patent: 10/5/1999) in view of Gupta et al. (US Patent No. 7,111,020, Filing Date: 3/26/2002).

Claim 2:

Regarding Claim 2, Sun discloses all the limitations above. However, Sun does disclose a method a method of calculating (column 6, lines 43-45, wherein two new rows with column primary key, i.e. CID, of 5 and 6 are added, resulting in snapshot, diagram 400 within Figure 4e, Sun). Sun is silent with respect to wherein the delta calculation module and a delta processing module calculates the plurality of delta values and the delta processing module directs a plurality of operators based on the operators.

On the other hand, Gupta discloses wherein a delta calculation module ("DCM") and a delta processing module ("DPM") in the module, wherein the DCM calculates the plurality of delta values and the DPM directs a plurality of operators based upon the plurality of delta values (column 2, lines 31-36; column 3, lines 51-57; and column 5, lines 61-67, Gupta).

It would have been obvious to one of the ordinary skill in the art at the time of the invention to incorporate a method for calculating and management of a database system by Gupta within Sun system to provide faster execution to reduce costly operations, and to improve the system performance.

Claim 4:

Regarding Claim 4, the combination of Sun in view of Gupta teaches wherein the DCM utilizes the plurality of group identifiers to combine the second plurality of data entries with the plurality of changes (Figures 1A and 1B, all features, Gupta).

Claim 6:

Regarding Claim 6, Sun in view of Gupta teaches a system for performing a pipelined refresh, the system comprising:

- a first materialized view derived at least partially from a base table (Figure 1B, diagram 142, Gupta);

- a refresh log having a plurality of entries, each of the plurality of entries corresponding to a change in the base table (Refer to claim 1, wherein this limitation has already been addressed, Sun), a second materialized view derived at least partially from the first materialized view (Figure 1B, diagram 144, Gupta);

- a refresh module that comprises:

 - a first delta calculation module that calculates a plurality of delta values that represents the changes to the first materialized view (column 2, lines 31-36, Gupta);

 - a first delta-processing module that applies the plurality of delta values to the first materialized view (column 2, lines 37-41, Gupta);

 - a delta adaptation module that receives the plurality of delta values from the first delta calculation module and calculates a plurality of changes to the second materialized view (column 3, lines 51-57, Gupta);

 - a second delta calculation module that obtains the plurality of changes to the second materialized view from the delta adaptation module (Figure 4B-1 and 2, and 4C-1 and 2, all features, wherein it illustrates a log used to track changes to base tables for an incremental refresh mechanism, Gupta); and

a second delta-processing module that applies the plurality of changes to the second materialized view from the second delta calculation module to the second materialized view (Figure 5A, diagrams 502, 504, and 506, Gupta).

Claim 7:

Regarding Claim 7, the combination of Sun in view of Gupta teaches wherein the plurality of entries in the refresh log correspond to a plurality of first materialized view entries in the first materialized view through a plurality of grouping identifiers that associate each of the plurality of entries with the plurality of first materialized view entries (Figure 2, all features, Gupta).

Claim 8:

Regarding Claim 8, Sun in view of Gupta teaches wherein a plurality of operators utilized by the first delta processing module to modify the first materialized view based upon the plurality of delta values (Figure 3A, all features, Gupta).

Claim 9:

Regarding Claim 9, the combination of Sun in view of Gupta teaches wherein the second delta calculation module is configured to calculate a plurality of second materialized view delta values from the plurality of changes and deliver the plurality of second materialized view delta values to the second delta processing module (Refer to claim 6, wherein these limitation are substantially the same/ or similar, Gupta).

Claim 10

Regarding Claim 10, the combination of Sun in view of Gupta teaches wherein the second delta-processing module is configured to utilize the plurality of second materialized

view delta values to apply the plurality of changes to the second materialized view (Refer to claim 6, wherein this limitation is substantially the same/or similar, Gupta).

Claim 18:

Regarding Claim 18, the combination of Sun in view of Gupta teaches wherein calculating the plurality of second materialized view delta values that represent the plurality of second materialized view changes to be applied to the second materialized view does not involve accessing a refresh log for the second materialized view (column 6, lines 27-31, Gupta).

Claim 19:

Regarding Claim 19, the combination of Sun in view of Gupta teaches wherein the method is performed in the recited order (column 14, lines 49-51, Gupta).

Claim 21:

Regarding Claim 21, the combination of Sun in view of Gupta teaches wherein each of the plurality of change entries comprises a group identifier (column 16, lines 12-16, wherein rows are grouped by customer_id and region, Gupta).

Claim 22:

Regarding Claim 22, the combination of Sun in view of Gupta teaches wherein the delta calculation module combines the plurality of change entries and a plurality of entries in the first materialized view to create the plurality of delta values (column 12, lines 3-10, Gupta).

(10) Response to Argument

1. I (Issue): did the Examiner err in concluding that claims 1, 3, 5, 11-17 and 20 were rejected under 35 U.S.C. 102(b) as being anticipated by Sun et al (US Patent No. 5,963,959).

- *In the first argument, the Appellant asserts, "The disclosure has been interpreted by the Examiner to correspond to the claimed module configured to access the refresh log and the first materialized view, as recited by the independent claims. This analysis incorrectly equates Sun's master table for Appellants' claimed refresh log. However, as appreciated by those skilled in the art, a master table is clearly not a refresh log. Therefore, the Sun reference does not teach the claimed module configured to access the refresh log and the first materialized view" (page 14).*

The Examiner does not agree with the Appellant's allegation because the refresh log was not equated to the master table, instead it was interpreted as "master log". In contrast, the Examiner pointed to the changes made to the master table and then recorded in the master log (Final Rejection, page 4, second paragraph).

- *In the second argument, the Appellant alleges that " the claimed module and its functionality pertaining to refresh operations, as recited by the claims are not and can not be equated with the user's ability to update a snap shot. Therefore, the Sun reference does not disclose or suggest a module adapted to perform a*

refresh operation on the first materialized view using the second plurality of data entries, as for example, recited by independent claim 1" (page 15).

The second argument has not been found persuasive. The Appellant alleges that equating refresh operation can not be equated to a snap shot, however there is no fact or explanation as to why such a comparison is improper. Furthermore, since the snapshot is at least a partial representation of a master view, it represents a first materialized view and a data entries corresponding to it (Figure 4(d), element 400, ZIP code). Furthermore the snapshot is updated as to include most current data entries with respect to the master view (Figures 4(d) and 4(e), note that new entries have been added to the element 400 from Figure 4(d) to Figure 4(e)). Consequently, the Examiner maintains that refreshing snapshot reads on the refresh operation.

- *In the third argument also on page 15, the Appellant alleges that "Sun's reference clearly does not teach a plurality of delta values, let alone calculating such delta values from a plurality of changes in a refresh log and second plurality of data values in a first materialized view".*

The Examiner does not agree with the Appellant's assertion. First of all, delta values similarly to the way of obtaining them is not defined by the Appellant neither in the claim, nor in the specification. It is well known in the art that delta values represent differences between two values or elements, and this is the meaning that the Examiner used in the interpretation of the claimed invention. Consequently, calculating plurality of

delta values corresponds to obtaining differences among portions of the corresponding entries. This is clearly illustrated in figures 2(c) and 2(d), wherein for the customer with ID 4, the zip code is changed from number 22090 to 22190 or figures 7(c) and 7(d) (i.e. snapshots representing master view). Furthermore, those changes are based on the master log (210) and the previous entries (200). Moreover the delta value has to be calculated in order to establish if the data in table 200 should be updated or it should remain the same (for instance is the newly submitted zip code any different from previously existing one).

- *In the fourth argument, the Appellant asserts that column 8, lines 52-67 "merely describes column entries of snapshot tables 700 and 710 but clearly does not teach a module configured to apply the plurality of delta values to the second plurality of data entries in the first materialized view" (page 17).*

The Appellant's fourth argument has not been found persuasive. As explained in the response to the previous argument, amending the zip code in the master table is actual step of applying delta values to the second data entries. First, the delta values are determined /calculated (i.e. the differences between the entries are noted) and then based on this determination, the data is modified or remains unchanged.

2. I(Issue): did the Examiner err in concluding that claims 2, 4, 6-10, 18,19, 21 and 22 stand rejected under 35 U.S.C. 103(a) as been unpatentable over Sun et al (US Patent No. 5,963,959) in view of Gupta (Patent No. 7,111,020).

- *In the fifth paragraph, the Appellant asserts that "Gupta teaches aggregate functions that are not used to obtain delta values, such as those derived from obtaining a difference between at least two values. The functions recited by the Gupta reference, such as SUM, COUNT and AVERAGE, are clearly not adapted to do so" (page 21).*

The Examiner agrees with the Appellant that mathematical operations such as SUM, COUNT and AVERAGE do not represent a difference between at least two values. However, the Examiner would like to note that Gupta's teaching does not only focus on calculating those aggregations but in addition, he also teaches refreshing data, not by exchanging all the data entries, but rather by obtaining differences between the content of the corresponding entries (column 3, lines 51-57). In order to accomplish such a task, Gupta has to calculate delta value indicating the difference between the data elements so that corresponding entries can be amended.

- *In the last argument also on page 21, the Appellant asserts that "Sun may disclose "changes" to base tables; there is no disclosure in Sun indicating what such changes represent. Particularly, Sun reference does not teach or suggest a first delta calculation module that calculates a plurality of delta values that represents the changes to the first materialized view, as specifically recited by independent claim 6. Based on the above analysis, it is apparent that Gupta does not teach, suggest or illustrate a first delta*

processing module that applies the plurality of delta values to the first materialized view”.

The Examiner does not agree with the Appellant's assertion. As it was explained above, the Appellant does not define delta value or the calculation required to obtain this value. In contrast, the Appellant associates the relationships and differences between particular data entries with delta values (i.e. differences). Consequently, the Examiner maintains that Sun teaches a step of determining differences between the data entries. For instance in figures 2(c) and 2(d), the difference between the zip codes for a customer with ID 4 is obtained. Furthermore, a person of ordinary skill in the art would note that updating the zip code involves calculating the delta value, wherein delta value describes the difference between the zip codes and therefore a need for an update of the information. Consequently, the customer ID remains the same while the second data entry is modified from 22090 to 22190.

(11) Related Proceeding(s) Appendix

For the above reasons, it is believed that the rejections should be sustained.

Respectfully submitted,



Angela M Lie

Conferees:

Angela M Lie (Examiner, AU 2163)

Tim Vo
(SPE, AU 2168)

~~Mohammed Ali (SPE, AU 2169)~~

Don Wong (SPE, AU 2163)



DON WONG
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100



TIM VO
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100